Semiparametric Regression for Discrete Time-to-Event Data

Moritz Berger [1] **and Matthias Schmid** [1]

[1] Department of Medical Biometry, Informatics and Epidemiology, Faculty of Medicine, University of Bonn, Germany

arXiv:1704.04087v1 [stat.AP] 13 Apr 2017

Address for correspondence: Moritz Berger, Department of Medical Biometry, Informatics and Epidemiology, Faculty of Medicine, University of Bonn, Siegmund-Freud-Straße 25, 53127 Bonn, Germany.

E-mail: Moritz.Berger@imbie.uni-bonn.de.

Phone: (+49) 228 287 14096.

Fax: (+49) 228 287 11324.

Abstract: Time-to-event models are a popular tool to analyse data where the outcome variable is the time to the occurrence of a specific event of interest. Here we focus on the analysis of time-to-event outcomes that are either intrisically discrete or grouped versions of continuous event times. In the literature, there exists a variety of regression methods for such data. This tutorial provides an introduction to how these models can be applied using open source statistical software. In particular, we consider semiparametric extensions comprising the use of smooth nonlinear functions and tree-based methods. All methods are illustrated by data on the duration of unemployment of U.S. citizens.

Key words: discrete time-to-event data; hazard models; semiparametric regression; survival analysis

1 Introduction

The objective of many statistical analyses is to model a duration time until a specifc event occurs. This is usually referred to as *time-to-event* or *survival analysis*. In biostatistics, for example, one often examines the time to death or the progression of a disease. In economics and the social sciences, popular examples include the modeling of the duration of unemployment or the time to retirement. Generally, in regression models for time-to-event data the event time itself is the response variable, and one wants to investigate the association of the response with several explanatory variables. Most often it is assumed in these analyses that the survival time is given by a random variable measured on a *continuous scale*. This case has been studied extensively in the literature, see, for example, Kalbfleisch and Prentice (2002) and Klein and Moeschberger (2003). However, in practice measurements of time are often discrete. Durations, for example, are often measured in days, years or months. Moreover, there are situations where the exact event time may not be known, but only an interval during which the event of interest took place.

Here we consider the application of regression models for *discrete* time-to-event data, which are characterized by an ordinal response variable taking the numbers $1, 2, \ldots$. These numbers either refer to a situation where event times are intrinsically discrete (such as the time to pregnancy, which in clinical applications is usually measured by

the number of menstrual cycles), or when continuous event times have been grouped. In the latter case, the numbers $t = 1, 2, \ldots$, refer to mutually exclusive time intervals $[0, a_1)$, $[a_1, a_2)$, $\ldots$, with fixed boundaries $a_1, a_2, \ldots$. Generally, a great advantage of discrete time-to-event models is that they can be viewed as regression models with binary response, giving rise, e.g., to the application of logistic regression or probit regression (Willett and Singer, 1993).

A comprehensive treatment of the statistical methodology for discrete time-to-event data has recently been given by Tutz and Schmid (2016). Similar to Gaussian regression, a large part of this methodology has been designed to estimate predictor-response relationships using a *linear* combination of the exaplanatory variables. In addition, Tutz and Schmid (2016) discuss several (less well known) approaches for *semiparametric* discrete time-to-event modeling. The aim of this tutorial is to provide an in-depth explanation of how these semiparametric models can be fitted and implemented using the `R` software for statistical computing (R Core Team, 2017). In particular, we will explain how smooth nonlinear functions and tree-based methods can be incorporated into discrete time-to-event models.

A frequently observed phenomenon in time-to-event analysis is censoring. Generally, a duration time is termed "censored" if its total length has not been fully observed. In this article we consider the most common type of *type-I* or *right censoring*, which means that the beginnings of the duration times are observed for all individuals in a study, whereas the respective ends are only observed for part of the individuals. Hence for some of the individuals it is only known that the event occured later than the observed time.

All models discussed in this article will be illustrated by means of a publicly available

Table 1: Summary statistics of the six explanatory variables used in the modeling of the U.S. unemployment data ($n = 3,210$).

Variable	Summary statistics					
	x_{min}	$x_{0.25}$	x_{med}	$\bar{x}$	$x_{0.75}$	x_{max}
`age`	20	27	34	35.45	43	61
`reprate`	0.06	0.39	0.50	0.45	0.52	2.05
`disrate`	0.01	0.05	0.10	0.11	0.15	1.02
`logwage`	2.70	5.29	5.68	5.69	6.05	7.60
`tenure`	0	0	2	4.11	5	40
`ui`	no: 1437 (44.8%)			yes: 1773 (55.2%)		

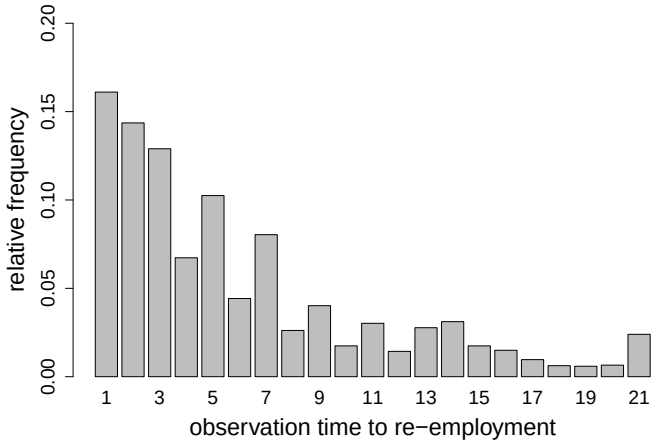

Figure 1: Observed time to re-employment (measured in two-week intervals) in the U.S. unemployment data. The median observation time in the data is 4, corresponding to a time period of 8 weeks.

data set on the duration of unemployment. The data comprise observations obtained from $n = 3,343$ U.S. citizens and were collected between 1986 and 1992 as part of the January Current Population Surveys Displaced Workers Supplements (DWS). The

original data set is available as part of the `R` add-on package `Ecdat` (Croissant, 2016). The response variable that will be considered here is the time to re-employment in any kind of job, which includes full-time, part-time or other kind of jobs. Due to the study design, the observed unemployment durations are discrete, as they were measured in two-week intervals. In this article, we will analyze the data over a period of 40 weeks comprising 21 possible event times $t = 1, 2, \ldots, 21$, where $t = 21$ refers to event times > 40 weeks. Explanatory variables that will be included in our analyses are the age in years (`age`), an indicator on whether an unemployment insurance claim was submitted (`ui`), the eligible replacement rate (`reprate`, defined by the weekly benefit amount divided by the amount of weekly earnings in the lost job), the eligible disregard rate (`disrate`, defined as the amount up to which recipients of unemployment insurance who accept part-time work can earn without any reduction in unemployment benefits divided by the weekly earnings in the lost job), the log weekly earnings in the lost job in \$ (`logwage`), and the tenure in the lost job in years (`tenure`). The summary statistics of the six explanatory variables are presented in Table 1. Due to missing values in the variables some observations were exluded from the data, arriving at a sample containing the data of $3,210$ citizens.

The observed times to re-employment are visualized in Figure 1. If an individual is still jobless at the end of the survey (i.e., after 40 weeks) or dropped out of the study before finding a job it is subject to right censoring. In this case its observation time corresponds to the *censoring time*, otherwise to the true time of re-employment.

The rest of this tutorial is organized as follows: Section 2 provides the basic theoretical framework and an introduction to parametric as well as semiparametric discrete time-to-event modeling. Details on model fitting and data preparation are given in

Section 3. Section 4 presents measures that are useful for assessing the goodness-of-fit of discrete time-to-event models. In Section 5 we illustrate the methods by presenting a detailed analysis of the U.S. unemployment data, showing how the various regression models can be applied by the use of `R`. Section 6 discusses additional aspects related to discrete time-to-event modeling and puts the methods considered in this article into perspective.

The `R` code to reproduce all the numerical results is provided as electronic supplement to this tutorial.

2 Notation and Basic Concepts

Given n observations, $i = 1, \ldots, n$, let in the following T_i denote the event time and C_i the censoring time of individual i. T_i and C_i are assumed to be independent random variables taking discrete values in $\{1, \ldots, k\}$. In addition one observes a vector of p explanatory variables $\boldsymbol{x}_i = (x_{i1}, \ldots, x_{ip})^\top$. For right censored data, the observation time is defined by $\tilde{T}_i = \min(T_i, C_i)$, i.e., $\tilde{T}_i$ corresponds to the true event time if $T_i < C_i$ and to the censoring time otherwise. If originally continuous data have been grouped, the discrete event times $1, \ldots, k$ refer to time intervals $[0, a_1), [a_1, a_2), \ldots, [a_{k-1}, \infty)$, where $T_i = t$ means that the event occured in time interval $[a_{t-1}, a_t)$. For example, in our application on unemployment durations, where time was measured in two-week intervals, $T_i = 3$ implies that re-employment of individual i took place between four and six weeks after the start of the study.

The main tool to model discrete time-to-event data is the *hazard function*, which captures the dynamics of the survival process at each time point. For a given vector

of explanatory variables $\boldsymbol{x}_i$, the hazard function is defined by

$$\lambda(t|\boldsymbol{x}_i) = P(T_i = t \,|\, T_i \geq t, \boldsymbol{x}_i), \quad t = 1, \ldots, k, \tag{2.1}$$

describing the conditional probability of an event at time t given that the individual survived until t. The corresponding *survival function* is given by

$$S(t|\boldsymbol{x}_i) = P(T_i > t|\boldsymbol{x}_i) = \prod_{s=1}^{t} (1 - \lambda(s|\boldsymbol{x}_i)), \quad t = 1, \ldots, k, \tag{2.2}$$

denoting the probability that an event occurs later than at time t, or, alternatively, the probability of surviving interval $[a_{t-1}, a_t)$.

An important consequence of the definition of the hazard function in (2.1) is that for a fixed time t, the hazard $\lambda(t|\boldsymbol{x}_i)$ describes a binary variable that distinguishes between the event taking place at time t or not, conditional on $T_i \geq t$. Therefore, a model for the discrete hazard function can be derived from regression modeling strategies for a binary response data.

2.1 Parametric Discrete Hazard Models

A general class of binary response models applied to the discrete hazard function is defined by

$$\lambda(t|\boldsymbol{x}_i) = h(\gamma_{0t} + \boldsymbol{x}_i^\top \boldsymbol{\gamma}), \tag{2.3}$$

where $h(\cdot)$ is a strictly monotone increasing distribution function. A common assumption is that the model contains a time-varying intercept and a set of covariate effects that are fixed over time. Hence, the linear predictor of the model, $\eta_{it} = \gamma_{0t} + \boldsymbol{x}_i^\top \boldsymbol{\gamma}$, comprises the intercepts γ_{0t}, $t = 1, \ldots, k-1$, and a vector of regression coefficients $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_p)^\top$ independent of t. Note that there is no intercept parameter for

$t = k$, as the hazard function in (2.1) is fully determined by $h(\cdot)$ and the coefficients $\gamma_{01}, \ldots, \gamma_{0,k-1}, \boldsymbol{\gamma}^\top$.

The most popular version of model (2.3) is the *logistic discrete hazard model* or *proportional continuation ratio model*, which is specified by the equation

$$\lambda(t|\boldsymbol{x}_i) = \frac{\exp(\gamma_{0t} + \boldsymbol{x}_i^\top \boldsymbol{\gamma})}{1 + \exp(\gamma_{0t} + \boldsymbol{x}_i^\top \boldsymbol{\gamma})}. \tag{2.4}$$

By definition, the proportional continuation ratio model uses the logistic distribution function for $h(\cdot)$. It can be shown that an alternative representation of the model is

$$\log\left(\frac{P(T_i = t|\boldsymbol{x}_i)}{P(T_i > t|\boldsymbol{x}_i)}\right) = \gamma_{0t} + \boldsymbol{x}_i^\top \boldsymbol{\gamma}. \tag{2.5}$$

The ratio $P(T_i = t|\boldsymbol{x}_i)/P(T_i > t|\boldsymbol{x}_i)$ compares the probability of an event at time t to the probability of an event later than t. It is also known as *continuation ratio*, see, for example, Agresti (2013). This representation of the model allows for an easy interpretation of the effects, see the application in Section 5.2.

Generally, the number of parameters in model (2.4) depends on the number of time points, as there is a seperate intercept for each t. The set of intercepts $\gamma_{01}, \ldots, \gamma_{0,k-1}$ defines the hazard that is always present for any given set of covariates. This hazard is ususally refered to as *baseline hazard*, and the intercepts γ_{0t} correspond to the log continuation ratio when all covariates are zero.

2.2 Semiparametric Extensions

The parametric model introduced in the previous section is linear in $\boldsymbol{\gamma}$, implying that each covariate has a linear effect on the transformed hazard. In practice, this linearity assumption may be too restrictive, as predictor-response relationships are

often characterized by nonlinear functional forms. Furthermore, it is assumed that the baseline hazard is represented by a separate intercept coefficient for each t. This can lead to numerical problems, especially when the number of time points (and hence the number of intercept parameters) is large relative to the sample size, implying that the event counts at some of these time points may become small. In the following we will consider popular semiparametric alternatives for the definition of η that address these issues. We first introduce *additive hazard models* and subsequently *tree-based methods*, which can also be embeeded in the framework of binary response models.

To avoid numerical problems in the estimation of the baseline hazard, it is often convenient to consider an additive model with predictor

$$\eta_{it} = f_0(t) + \boldsymbol{x}_i^\top \boldsymbol{\gamma}, \tag{2.6}$$

where $f_0(t)$ is a smooth (possibly nonlinear) function of time. By relating the values of the baseline hazard at neighboring time points via $f_0(t)$, the number of parameters involved in model fitting effectively reduces, and low event counts at some time points become less problematic. A common way to specify the smooth function in t is to use splines, which are represented by a weighted sum of m basis functions. One possible representation of $f_0(t)$ is by B-spline basis functions. These are polynomials of fixed degree d differing from zero in $d+1$ adjacing intervals. For a comprehensive introduction to B-splines, see De Boor (1978). Very flexible spline functions can be obtained by choosing a relatively large number of basis functions m and at the same time using a penalty term to prevent estimates becoming too rough ("wiggly"). This approach, on which we will focus in this article, is called *P-splines* and was first proposed by Eilers and Marx (1996).

An extension of the semiparametric model (2.6) that weakens the linearity assumption

on the effects of the covariates is given by the additive model

$$\eta_{it} = f_0(t) + \sum_{j=1}^{p} f_j(x_{ij}), \tag{2.7}$$

where the $f_j(x_j)$ are unknown smooth functions. That is, the effects of the covariates (or subsets of the covariates) are determined by smooth, possibly nonlinear, functions. A common approach is again to use P-splines and to expand each function seperately by a weighted sum of B-spline basis functions depending on the covariates.

In the semiparametric models with predictors (2.6) and (2.7) it is assumed that the predictor is given by an additive function of time and a linear (or additive) function of the covariates. Although these models are very flexible, they may not capture the structure of the data very well if interactions between covariates are present. For example, it is quite conceivable that the effect of a covariate on the hazard depends on the values of a second covariate, implying the presence of an interaction between the two covariates. The problem when incorporating interactions in parametric or additive models is that the relevant interactions have to be known and specified before model fitting. Furthermore, parametric and additive models are hard to handle if the interaction terms involve more than two covariates. An alternative regression approach that addresses these problems is *recursive partitioning*, which is also known as *tree modeling*. The most popular tree method is *classification and regression trees* (CART), as proposed and described in detail by Breiman et al. (1984). The basic CART method is conceptually very simple: The covariate space is partitioned recursively into a set of rectangles, and in each rectangle a simple model (for example, a covariate-free model) is fitted. A user-friendly introduction to the basic concepts of tree modeling is found in Hastie et al. (2009). Recently Schmid et al. (2016) proposed a recursive partitioning method that is specifically designed to model discrete

time-to-event data. The main principle is to fit a discrete hazard model of the form

$$\lambda(t|\boldsymbol{x}_i) = f(t, \boldsymbol{x}_i), \tag{2.8}$$

where $f(t, \boldsymbol{x})$ is represented by a classification tree with binary outcome. Each split of this tree is determined by either t (treated as an ordinal variable) or one of the covariates. As a result, each terminal node of the tree refers to an estimate of the hazard function for a specific covariate combination and a specific time interval $[t_1, t_2] \subset [1, k]$. For details on the calculation of the estimates, see Section 5.3.

3 Estimation and Data Preparation for Additive Hazard Models

To derive the log-likelihood function for discrete hazard models it is useful to introduce a binary variable indicating whether the target event was observed or not:

$$\Delta_i = \begin{cases} 1, & \text{if } T_i \leq C_i, \\ 0, & \text{if } T_i > C_i. \end{cases} \tag{3.1}$$

Thus, Δ_i becomes 1 if the exact true time is observed; otherwise, $\Delta_i = 0$. In the case where continuous time-to-event data are grouped, $\Delta_i = 1$ and $\tilde{T}_i = t$ implies an event in interval $[a_{t-1}, a_t)$ and $T_i = \tilde{T}_i = t$. Similarly, $\Delta_i = 0$ and $\tilde{T}_i = t$ implies $C_i = \tilde{T}_i = t$ and survival beyond a_t, i.e., $T_i > \tilde{T}_i = t$.

Note that when continuous time-to-event data are grouped or rounded, additional assumptions are implicitly imposed on the censoring mechanism. To see this, consider the case where both the continuous event time $T_{\text{cont},i}$ and the continuous censoring

time $C_{\text{cont},i}$ are within the same interval, say, $[a_{t-1}, a_t)$. Then, by definition, $T_i = C_i = t$, $\tilde{T}_i = t$, and $\Delta_i = 1$, leading to the usual interpretation that an event was observed in interval $[a_{t-1}, a_t)$. At the same time, however, this interpretation implicitly assumes $T_{\text{cont},i} \leq C_{\text{cont},i}$, i.e., the continuous event time $T_{\text{cont},i} \in [a_{t-1}, a_t)$ is not allowed to be larger than the continuous censoring time $C_{\text{cont},i} \in [a_{t-1}, a_t)$. Without this assumption, the scenario where both $T_{\text{cont},i}, C_{\text{cont},i} \in [a_{t-1}, a_t)$ and $T_{\text{cont},i} > C_{\text{cont},i}$ would result in $\tilde{T}_i = t$ and $\Delta_i = 1$ but *no* observed event in $[a_{t-1}, a_t)$, implying that the usual interpretation of $\tilde{T}_i = t$, $\Delta_i = 1$ would no longer be appropriate. This assumption on the nature of the censoring mechanism is often referred to as "censoring at the end of the interval".

With data $(\tilde{T}_i, \Delta_i, \boldsymbol{x}_i)$, $i = 1, \ldots, n$, the contribution of the i-th observation to the likelihood function is given by

$$L_i = P(T_i = \tilde{T}_i)^{\Delta_i} \, P(T_i > \tilde{T}_i)^{1-\Delta_i} \, P(C_i \geq \tilde{T}_i)^{\Delta_i} \, P(C_i = \tilde{T}_i)^{1-\Delta_i}. \tag{3.2}$$

A crucial assumption that is usually made to simplify the likelihood function (3.2) is that the censoring process does not depend on the parameters determining the event times T_i. A consequence of this assumption is that the terms involving the censoring times can be ignored in the maximization of the likelihood function for the time-to-event process. Omitting the terms involving C_i in (3.2) and inserting the definitions of the hazard function (2.1) and the survival function (2.2), one obtains (expect for some constants)

$$L_i \propto \lambda(\tilde{T}_i|\boldsymbol{x}_i)^{\Delta_i}(1 - \lambda(\tilde{T}_i|\boldsymbol{x}_i))^{1-\Delta_i} \prod_{j=1}^{\tilde{T}_i - 1} (1 - \lambda(j|\boldsymbol{x}_i)). \tag{3.3}$$

Note that, by definition, one always obtains $\Delta_i = 1$ and $\lambda(\tilde{T}_i|\boldsymbol{x}_i) = 1$ if $\tilde{T}_i$ is equal to the last time point k. For maximum likelihood estimation, it is therefore convenient

to re-code observations with $\tilde{T}_i = k$ as follows:

$$\tilde{T}_i = k,\ \Delta_i = 1,\ \boldsymbol{x}_i \ \longmapsto\ \tilde{T}_i = k-1,\ \Delta_i = 0,\ \boldsymbol{x}_i\ , \tag{3.4}$$

making use of the fact that the value of the likelihood contribution in (3.3) will not be altered by this transformation.

With some algebra it can be shown that the likelihood function (3.3) is equal to the likelihood of a binary response model with outcome variables

$$(y_{i1}, \ldots, y_{i\tilde{T}_i}) = \begin{cases} (0, \ldots, 0, 1), & \text{if } \Delta_i = 1 \\ (0, \ldots, 0, 0), & \text{if } \Delta_i = 0. \end{cases} \tag{3.5}$$

For individuals where the exact event time is observed one defines the observation vector $(0, \ldots, 0, 1)$ of length $\tilde{T}_i$. For censored individuals the observation vector contains only zeros. According to this definition one has $\tilde{T}_i$ binary observations for each individual i, resulting in a total of $\tilde{T}_1 + \ldots + \tilde{T}_n$ observations. Using these definitions, the log-likelihood of the proportional continuation ratio model becomes

$$l \ \propto \sum_{i=1}^{n} \sum_{s=1}^{\tilde{T}_i} y_{is} \log(\lambda(s|\boldsymbol{x}_i)) + (1 - y_{is}) \log(1 - \lambda(s|\boldsymbol{x}_i)). \tag{3.6}$$

The main advantage of this representation of the log-likelihood is that it allows to use software for fitting binary response models. For example, it follows from (3.6) that fitting a continuation ratio model is equivalent to fitting a logistic regression model with predictor (2.6) or (2.7). In this model the values of the binary responses y_{is} can be interpreted as binary decisions for the transition from interval $[a_{s-1}, a_s)$ to $[a_s, a_{s+1})$. For instance, in the application on unemployment duration one observes $y_{is} = 0$ for each two-week interval as long as the individual i is not re-employed yet.

The models with smooth components (2.6) and (2.7) can be fitted by maximizing a

penalized likelihood of the form

$$\ell_p = \ell - \delta J, \tag{3.7}$$

where $\delta \in \mathbb{R}^+$ is a penalty parameter and $J \in \mathbb{R}^+$ is the penalty term mentioned in Section 2.2 putting restrictions on the weights of the B-spline basis functions and preventing estimates from becoming too rough. When using P-splines, J is a difference penalty on adjacent B-spline coefficients, see, e.g., Eilers and Marx (1996) for details. A common procedure is to use cubic B-splines ($d = 3$) with second order differences. The degree of smoothness is determined by the tuning parameter δ. The larger the value of δ, the smoother is the resulting function, and vice versa. When several smooth functions are included in the model, one uses a difference penalty for each spline effect, based on the differences of adjacent B-spline coefficients for the corresponding covariate. The smoothness of the individual spline estimates is determined by seperate penalty parameters δ_j.

Before fitting proportional continuation ratio models with software for binary outcome data, one has to generate the required binary observations presented in (3.5). This is done by the generation of an *augmented data matrix*. For the setup of the matrix one has to distinguish between censored and non-censored individuals. For an individual whose event was observed ($\Delta_i = 1$) at time $\tilde{T}_i$ the augmented data matrix is given by

$$\begin{pmatrix} 0 & 1 & x_{i1} & \dots & x_{ip} \\ 0 & 2 & x_{i1} & \dots & x_{ip} \\ 0 & 3 & x_{i1} & \dots & x_{ip} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & \tilde{T}_i & x_{i1} & \dots & x_{ip} \end{pmatrix}. \tag{3.8}$$

For an individual that is censored ($\Delta_i = 0$) at time $\tilde{T}_i$ the augmented data matrix is

given by

$$\begin{pmatrix} 0 & 1 & x_{i1} & \dots & x_{ip} \\ 0 & 2 & x_{i1} & \dots & x_{ip} \\ 0 & 3 & x_{i1} & \dots & x_{ip} \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & \tilde{T}_i & x_{i1} & \dots & x_{ip} \end{pmatrix}. \tag{3.9}$$

The first column in the augumented data matrices corresponds to the binary responses $y_{i1}, \dots, y_{i\tilde{T}_i}$. The second column is the time interval running from 1 to $\tilde{T}_i$. When fitting a model with fixed intercept parameters γ_{0t} this column has necessarily to be coded as a nominal factor, e.g., via dummy variables. The remaining part of the data contains the covariates. When the covariates are constant over time, the values in each row of columns 3 to $(p + 2)$ are the same. This is also the case in the U.S. unemployment data. Otherwise, when there are time-varying covariates, the observed time series are entered in the respective columns of the augmented data matrices. For each individual the augumented data matrix has $\tilde{T}_i$ rows, and the whole data matrix, which is obtained by "glueing" the individual augmented matrices together, has $\sum_{i=1}^{n} \tilde{T}_i$ rows.

In `R` the augumented data matrix can be generated by applying the function `dataLong()` in the `R` package **discSurv** (Welchowski and Schmid, 2017). The general interface of the function is

```
> dataLong(dataSet, timeColumn, censColumn, timeAsFactor = TRUE)
```

The function requires the original data of class `data.frame` in "non-augumented" short format (argument `dataSet`), the column name of the observed discrete event times (argument `timeColumn`) and the column name of the binary event indica-

tor as defined in equation (3.1) (argument `censColumn`). The variable required by `timeColumn` can either be numeric or coded as an ordinal or nominal factor. If `timeAsFactor = TRUE` the time column in the augumented data matrix will be returned as a nominal factor. The variable required by `censColumn` can either be a numerically coded 0/1 vector or a labeled factor variable. Note that `dataLong()` assumes that the covariates are constant over time. If this is not the case the function `dataLongTimeDep()` should be used instead to generate the augumented data matrix.

The augumented data matrix returned by `dataLong()` contains the binary responses as defined in equation (3.5) in the form of a numerically coded 0/1 vector named `y`. Further details on the output are given by the application in Section 5.1.

4 Goodness-of-Fit Measures

In this tutorial we consider two diagnostic tools that are useful to investigate discrete hazard models in terms of their goodness of fit. First, one can generate a *calibration plot*. The idea is to compare the estimated hazards $\hat{\lambda}(t|\boldsymbol{x}_i), i = 1, \ldots, n,\ t = 1, \ldots, \tilde{T}_i$, of the model to the relative frequencies of observed events ($y_{it} = 1$) in predefined subsets of the augmented set of observations. More specifically, one splits the data into subsets $D_k,\ k = 1, \ldots, K$, defined by the percentiles of the estimated hazards. Common choices for K are $K = 10$ or $K = 20$. Then the relative frequency of observed events ("empirical hazard") is calculated in each subset by

$$\sum_{i,t:\hat{\lambda}(t|\boldsymbol{x}_i)\in D_k} \frac{y_{it}}{|D_k|}\,, \tag{4.1}$$

where $|D_k|$ corresponds to the number of observations in subset D_k. If the fit of the model is satisfactory the empirical hazard measure in (4.1) should be close to the average of the estimated hazards in D_k for all k. An example of a calibration plot is shown in Figure 3 in the application.

Second, we consider *martingale residuals*, which allow for assessing the importance of single covariates x_j. The idea of the martingale residuals is to compare for each individual the observed number of events with the expected number of events up to $\tilde{T}_i$. Using the binary response variables $y_{i1}, \ldots, y_{i\tilde{T}_i}$ the residuals are defined as

$$r_i = \sum_{t=1}^{\tilde{T}_i} (y_{it} - \hat{\lambda}(t|\boldsymbol{x}_i)), \quad i = 1, \ldots, n. \tag{4.2}$$

For a well fitting model that includes all relevant predictors, the difference between y_{it} and $\hat{\lambda}(t|\boldsymbol{x}_i)$ should be "random" and therefore uncorrelated with the covariate values. To assess the importance of a covariate graphically one can plot the residuals against the covariate values. Martingale residuals can be computed by the function `martingaleResid()` contained in the **discSurv** package. An example is shown in Figure 3 in the application.

5 Application: Duration of Unemployment

In the following discrete hazard regression modeling is illustrated by means of a step-by-step analysis of the U.S. unemployment data. Throughout this section we use the logistic link function, i.e., we consider the fitting of a proportional continuation ratio model.

5.1 Preprocessing of the Data

To fit a logistic discrete hazard model of the form (2.4), the original data matrix first has to be transformed to an augmented data matrix, as described above. The data set `UnempDur`, which (after application of the pre-processing steps outlined in the Introduction) is a slightly modified version of the data frame available in the `R` package `Ecdat`, has the following form:

```
> head(UnempDur)
```

```
  spell age  ui reprate disrate logwage tenure status
1     5  41  no   0.179   0.045 6.89568      3      1
2    13  30 yes   0.520   0.130 5.28827      6      1
4     3  26 yes   0.448   0.112 5.97889      3      1
5     9  22 yes   0.320   0.080 6.31536      0      1
6    11  43 yes   0.187   0.047 6.85435      9      0
8     3  32  no   0.373   0.093 6.16121      0      1
```

The first column named `spell` is the observed time to re-employment of indiviudal i and contains the values of $\tilde{T}_i$, $i = 1, \ldots, n$. As mentioned above, these values correspond to the lengths of the spells (measured in two week intervals), whose distribution is displayed in Figure 1. The last column named `status` indicates whether the exact event time of invidiual i has been observed (`status` $= 1$) or if the individual is subject to right censoring (`status` $= 0$); it corresponds to the random variable Δ_i defined in equation (3.1). Summarizing the `status` column yields a censoring rate of 0.391:

```
> table(UnempDur$status)/nrow(UnempDur)

        0         1
0.3909657 0.6090343
```

Columns two to five of the data frame `UnempDur` contain the explanatory variables described in Table 1. All covariates are constant over the time of the survey.

When using `dataLong()` to obtain the augumented data matrix one has to pass the column names `spell` and `status` to the arguments `timeColumn` and `censColumn`, respectively:

```
> library(discSurv)
> UnempDurLong <- dataLong(UnempDur, timeColumn = "spell",
+                          censColumn = "status")
```

The augmented data matrix `UnempDurlong` has ten columns with the following names:

```
> names(UnempDurLong)

 [1] "obj"     "timeInt" "y"       "spell"   "age"     "ui"
 [7] "reprate" "disrate" "logwage" "tenure"  "status"
```

The new columns are `obj`, which is an identifier of the individuals, `timeInt`, which contains the discrete time values (i.e., the second column of the augmented matrices in (3.8) and (3.9), stored as a nominal factor) and `y`, which contains the binary response variables $y_{i1}, \ldots, y_{i\tilde{T}_i} \in \{0, 1\}$. The head of the augmented data matrix is given by

```
> UnempDurLong[UnempDurLong$obj==1, ]
```

```
    obj timeInt y spell age ui reprate disrate logwage tenure status
1     1       1 0     5  41 no   0.179   0.045 6.89568      3      1
1.1   1       2 0     5  41 no   0.179   0.045 6.89568      3      1
1.2   1       3 0     5  41 no   0.179   0.045 6.89568      3      1
1.3   1       4 0     5  41 no   0.179   0.045 6.89568      3      1
1.4   1       5 1     5  41 no   0.179   0.045 6.89568      3      1
```

showing that the first individual (`obj` $= 1$) had an event after ten weeks (`spell` $= 5$ and `status` $= 1$). Accordingly, the augumented data matrix for the first individual has five rows, where each row corresponds to one time interval (`timeInt` $= 1, \ldots, 5$). The corresponding vector of responses is `y`$= (0, 0, 0, 0, 1)$. The values of the covariates remain constant over time and are therefore the same in each row.

As a second example, consider the augmented data matrix of the twelfth invidiual (`obj`$= 12$). This individual is censored after 6 weeks (`spell`$= 3$ and `status`$= 0$), and hence the corresponding data matrix has three rows with response `y`$= (0, 0, 0)$:

```
> UnempDurLong[UnempDurLong$obj==12, ]
```

```
     obj timeInt y spell age  ui reprate disrate logwage tenure status
14    12       1 0     3  40 yes    0.52    0.13 4.95583      0      0
14.1  12       2 0     3  40 yes    0.52    0.13 4.95583      0      0
14.2  12       3 0     3  40 yes    0.52    0.13 4.95583      0      0
```

5.2 Regression Modeling

A parametric proportional continuation ratio model with a linear predictor is estimated in R by passing the augumented data matrix `UnempDurLong` to `glm()` with the usual specifications:

```
> model1 <- glm(formula = y ~ timeInt - 1 +
+               age + reprate  + disrate + logwage + tenure + ui,
+               data = UnempDurLong, family = binomial(link = "logit"))
```

The left hand side of the `formula` argument contains the binary response vector `y`. In addition to the names of the six covariates, the right hand side contains the discrete time variable as a nominal factor without intercept (`timeInt - 1`). The `family` argument `binomial(link = "logit")` is the same as for "usual" logistic regression models with binary outcome.

The more complex model with smooth baseline hazard (equation (2.6)) is estimated by use of the R package **mgcv** (Wood, 2017). A detailed introduction to the estimation procedures is found in Wood (2011). The corresponding function `gam()` essentialy has the same interface as `glm()`:

```
> library("mgcv")
> UnempDurLong$timeIntNum <- as.numeric(UnempDurLong$timeInt)
> model2 <- gam(formula = y ~ s(timeIntNum, bs = "ps", k = 5, m = 2) +
+               age + reprate  + disrate + logwage + tenure + ui,
+               data = UnempDurLong, family = binomial(link = "logit"))
```

Before passing the nominal factor `timeInt` to `gam()`, it has to be transformed to a continuous variable (`timeIntNum`). Then a smooth baseline hazard is specified by use

of the function `s()` on the right hand side of the model formula. Required arguments are the type of spline smoother `bs`, the dimension of the basis `k`, which determines the number of basis functions, and the order of the penalty `m`. Here we use P-splines with cubic basis functions (`bs = "ps"`) and a second order difference penalty (`m` $= 2$). The chosen dimension (`k` $= 5$) results in nine cubic basis functions. Based on these specifications the optimal smoothing parameter δ, see equation (3.7), is computed by generalized cross-validation, see Wood (2006). Its estimate is stored in the argument `sp`. In case of the U.S. unemployment data, `sp` is estimated as:

```
> model2$sp


s(timeIntNum)
   0.08562284
```

Generally, the **mgcv** package implements a large variety of alternative spline estimators and methods for smoothing parameter optimization. In principle, all of these methods may used for discrete hazard modeling, in the same way as they would be used in logistic regression (or, more generally, in additive models with a binary response).

The estimated baseline hazards of `model1` and `model2` are shown in Figure 2. The discrete baseline hazard obtained for `model1` is visualized by black dots. From these estimates a reasonable interpretation is hard to derive. On the other hand, the smooth baseline hazard obtained for `model2` visualized by grey squares is more meaningful. It is seen that the conditional probability of re-employment decreases until week 20 and subsequently increases up to week 32 before it diminishes again. The reason for this might be that in many U.S. states workers are eligible for up to 26 weeks of benefits

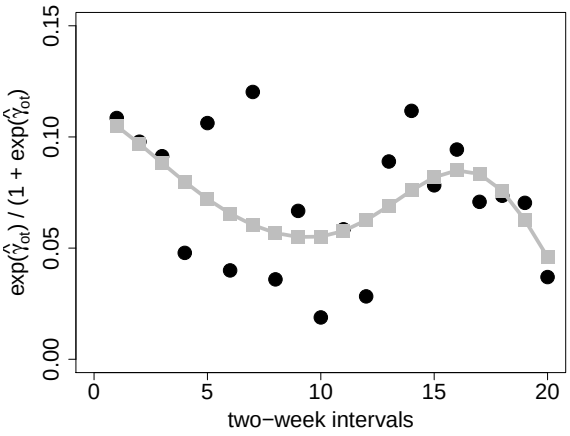

Figure 2: Analysis of the U.S. unemployment data. The figure shows the estimated discrete baseline hazard of `model1` (black dots) and the smooth baseline hazard of `model2` (grey squares).

Table 2: Analysis of the U.S. unemployment data. The table contains coefficient estimates (coef), estimated standard errors (se) and p-values of the covariate effects obtained for `model1` and `model2` (bh = baseline hazard).

	model1 (discrete bh)			**model2 (smooth bh)**		
	coef	se	p-value	coef	se	p-value
age	-0.012	0.003	0.000	-0.012	0.003	0.000
reprate	0.285	0.342	0.406	0.301	0.338	0.373
disrate	-0.764	0.383	0.046	-0.755	0.379	0.047
logwage	0.231	0.072	0.001	0.236	0.071	0.001
tenure	-0.005	0.005	0.280	-0.006	0.005	0.266
ui	-1.151	0.052	0.000	-1.175	0.051	0.000

from the state-funded unemployment compensation program.

Table 2 shows the estimates of the coefficients $\boldsymbol{\gamma}$, the corresponding estimated stan-

dard errors, and the p-values of the covariate effects obtained for `model1` and `model2`. Apart from the eligible replacement rate (`reprate`) and the tenure in the lost job (`tenure`), the covariates are significantly associated with the time to re-employment. According to the signs of the estimates of both models, the chance of getting re-employed decreases with increasing values of age, increasing disregard rate and with the filing of an unemployment claim. On the other hand, the higher the earnings in the lost job the better the chance of re-employment. Table 2 also shows that the differences in coefficient estimates between the two models are small. By use of equation (2.5) the effects can be interpretated in an easy way: For example, let us compare citizens who submitted an unemployment insurance claim (`ui` $= 1$) to those who did not (`ui` $= 0$). Based on the estimate of `model1` ($\hat{\gamma}_{ui} = -1.151$), one obtains that the probability of re-employment at time t, compared to the probability of re-employment later than t, decreases for citizens who submitted an unemployment insurance claim by the factor $\exp(\hat{\gamma}_{ui}) = 0.316$. The chance of re-emplyoment is therefore much smaller in this group. One might speculate that due to benefits from the state-funded unemployment the motivation to search for a new job is lower.

The goodness-of-fit measures for `model1` are presented in Figure 3. The left panel shows the calibration plot (average fitted hazards against the relative frequencies of events). It is seen that the values do not deviate too much from the 45 degree line, indicating an acceptable model fit. The right panel shows the martingale residuals defined in (4.2), against the values of the covariate `age` for `model1` without age. The black line corresponds to the estimated trend obtained by a local polynomial regression using the `R` function `loess()`. The functional form of the trend line (compared to the zero line) shows a non-linear effect on the martingale residuals. This indicates that the covariate `age` is an influential variable with a non-linear effect on the

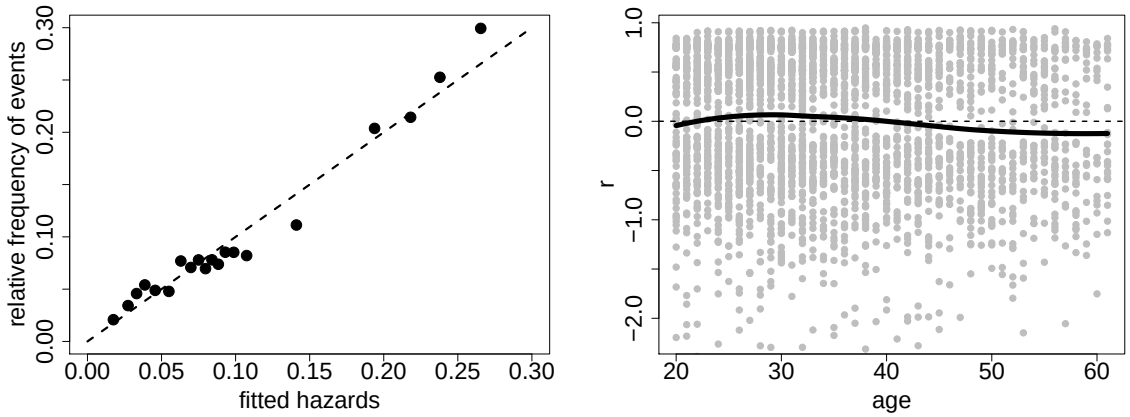

Figure 3: Analysis of the U.S. unemployment data. The two panels show the calibration plot (left) for `model1` and the martingale residuals against the values of age (right) obtained for `model1` without age. The trend line was obtained by local polynomial regression.

response.

Therefore, as a possible extension, we consider a model where the baseline hazard as well as the covariate age are both modeled as smooth P-spline functions:

```
> model3 <- gam(formula = y ~ s(timeIntNum, bs = "ps", k = 5, m = 2) +
+               s(age, bs = "ps", k = 25, m = 2) +
+               reprate  + disrate + logwage + tenure + ui,
+               data = UnempDurLong, family = binomial(link = "logit"))
```

As seen from the `R` code, the estimation of the smooth function of covariate age in `model3` is based on 29 cubic basis functions (dimension `k` = 25). The estimated penalty parameter $\hat{\delta}$ for age, stored in `sp`, is:

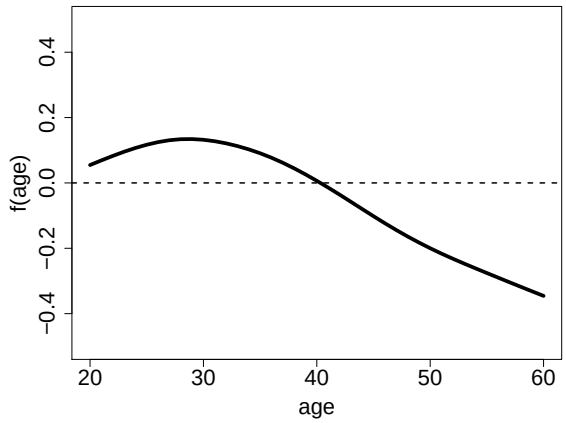

Figure 4: Analysis of the U.S. unemployment data. The plot shows the estimated P-spline function for the covariate age in `model3`.

```
> model3$sp["s(age)"]
```

```
 s(age)
56860.2
```

From the resulting function shown in Figure 4 it is seen that the association between the time to re-employment and age is definitely not linear. Its form is very similar to the loess trend shown in Figure 3. The value on the y-axis of the figure corresponds to the contribution of age to the predictor η_{it} of the model. The chance of re-employment has a peak between 20 years and 30 years and subsequently decreases.

5.3 Tree-Based Modeling

Finally, we fit a recursive partitioning model of the form (2.8). Again we consider a procedure that is based on the augmented data matrix with binary outcomes $y_{i1}, \ldots, y_{i\tilde{T}_i}$.

When growing trees one has to take two main decisions: Firstly, one has to choose an appropriate criterion for performing the splits. Criteria that have already been used in the early days of tree construction are impurity measures. For discrete surival trees a natural measure of node impurity is the *Brier score*, which evaluates the average squared difference between the binary outcome values y_{it} and the respective hazard estimate $\hat{\lambda}(t|\boldsymbol{x}_i)$ in each node, see Schmid et al. (2016). It can be shown that using the Brier score is equivalent to the traditional Gini impurity measure. For a single node m the Gini impurity is given by

$$G_m = 2\,\pi_m\,(1 - \pi_m), \tag{5.1}$$

where π_m is the proportion of ones in node m, see Breiman (1996). This equivalence implies that the traditional CART algorithm based on the Gini criterion can be used for the construction of the tree. The latter is done by using the function `rpart()` of the eponymous `R` package **rpart** (Therneau et al., 2015).

Secondly, one has to determine the optimal size of the tree. For the discrete survival tree an appropriate tuning parameter controlling tree size is the minimal number of observations in each terminal node ("minimal node size"). Optimizing this number avoids overfitting, as the number of terminal nodes is prevented from becoming too large and, at the same time, the node sizes are prevented from becoming too small. Accordingly, splitting is stopped when further splitting in any of the current nodes would result in an additional node containing less observations than the minimal node size. Given a sequence of tree estimates depending on the minimal node size, the optimal tree (i.e., the tree with "optimal" minimal node size) is determined by either minimization of an information criterion (such as AIC or BIC, see below) or maximization of the predictive log-likelihood. The latter strategy means to repeat-

edly draw subsamples from the original non-augmented data (for example, by cross-validation, bootstrapping or subsampling without replacement), and to calculate the log-likelihood for the omitted observations. One determines the optimal tree as the one for which the predictive log-likelihood (averaged across the subsamples) becomes maximal. The `R` function `survivalTree()` automatically generates the augumented data matrix by `dataLong()`, estimates the discrete survival tree by `rpart()` und returns the optimal one according to the specified performance criterion. The function is part of the electronic supplement of this article.

Once the optimal minimal node size has been determined, the estimate of $\lambda(t|\boldsymbol{x}_i)$ is given by the relative frequency of events (proportion of ones) in each node, possibly after applying some sort of correction procedure like the Laplace correction (see below).

To fit a tree model to the U.S. unemployment data, we call the `survivalTree` function using the following arguments:

```
> source("survivalTree.R")
> model4 <- survivalTree(formula = y ~ timeInt + age +
+                        reprate + disrate + logwage + tenure + ui,
+                        data = UnempDur, tuning = "BIC",
+                        timeColumn = "spell", censColumn = "status",
+                        minimal_ns = seq(100, 1500, by = 10),
+                        trace = TRUE)
```

The formula required for the tree model is analogous to the one specified for a model with linear predictor. Note that internally the time variable `timeInt` is coded as a numeric vector. This is in analogy to `model2` and `model3` with smooth baseline hazard.

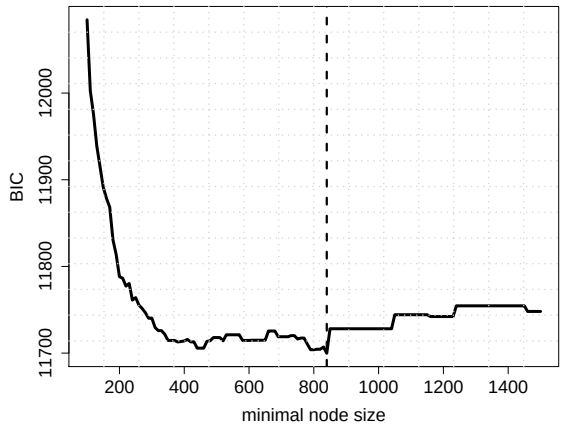

Figure 5: Analysis of the U.S. unemployment data. The plot shows the BIC values for the sequence of survival trees that was obtained by fitting `model4` with minimal node sizes ranging from 100 to 1500. The minimal BIC value (obtained for node size 840) is marked by the vertical dashed line.

The original data frame `UnempDur` (in non-augmented format) is passed to the `data` argument. In addition one has to specify the `timeColumn` and `CensColumn` arguments used in `dataLong()`. The performance criterion is specified by the argument `tuning`. For tuning we use the Bayesian information criterion (BIC) defined by

$$\text{BIC} := -2\,l + \log(\tilde{n})\,n_s\,,$$

where l is the log-likelihood (3.6), $\tilde{n}$ is the number of rows of the augumented data matrix, and n_s denotes the number of splits as a measure of the complexity of the tree. Other possible arguments for tuning are `"AIC"` (Akaike's information criterion) and `"ll"` (predictive log-likelihood method). When using `"ll"`, `survivalTree()` performs a five-fold cross-validation based on subsamples without replacement stratified by `spell`. The `survivalTree` function searches for the best model among the sequence of models with minimal node sizes `minimal_ns`. If `minimal_ns` is not specified, the sequence of minimal node sizes is set to $1, \ldots, \lfloor \tilde{n}/2 \rfloor$.

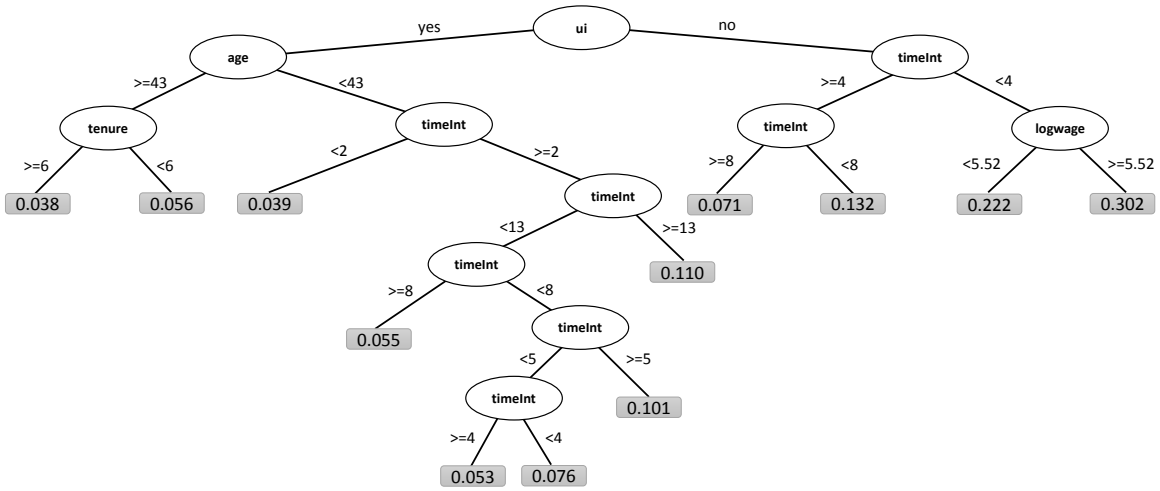

Figure 6: Analysis of the U.S. unemployment data. The graph visualizes the survival tree obtained from fitting `model4` with BIC-optimal minimal node size 840. The numbers at the terminal nodes refer to the estimated hazards. All estimated hazards were additionally post-processed by application of the *Laplace correction*, which was suggested by Ferri et al. (2003) to correct for estimates near the boundaries 0 and 1 in nodes with very few observations. The Laplace correction is automatically performed by `survivalTree()`.

The BICs obtained for `model4` with minimal node sizes $100, \ldots, 1500$ (in steps of 10) are shown in Figure 5. If an increase of the minimal node size does not change the number of splits and therefore does not influence the resulting tree, the BIC remains the same. This is the case, for example, between minimal node sizes 900 and 1000. According to the BIC, the optimal tree model has minimal node size 840, marked by the dashed line in Figure 5. This results in a tree with eleven splits or twelve terminal nodes. The estimated tree is shown in Figure 6.

The most important covariate, which was chosen in the first split of the tree, is `ui`. As already derived from the parametric models, the submission of an unemployment

insurance claim (`ui` = "yes") has a negative effect on the "chance" of re-employment. Within the group of citizens who submitted an insurance claim, the chance is lowest for citizens aged 43 years or older and with a tenure in the lost job of at least 6 years (leftmost node in Figure 6). For citizens younger than 43 years all further splits are performed with regard to the discrete time variable (`timeInt`). This confirms the results from Figure 2 in that the chance of re-employment is highly time-dependent. With a hazard estimate of 0.110 after 26 weeks (`timeInt` >= 13) of unemployment in this group, the tree estimate also reflects the marked increase in the chance of re-employment already seen in Figure 2 after 20 weeks. The best opportunities of re-employment are observed for citizens without an unemployment insurance claim, within the first six weeks (`timeInt` < 4) of unemployment, and with log weekly earnings of at least 5.52$ (rightmost node in Figure 6). For this subgroup the estimated hazard rate is 0.302. The two covariates `reprate` and `disrate` were not selected in any of the splits and are therefore exluded from the model. This is in contrast to `model1` and `model2` (see Table 2), where `disrate` showed a significant effect on the hazard.

6 Concluding Remarks

In this tutorial we have described a basic set of tools to fit semiparametric regression models with a discrete time-to-event outcome. All presented models are very general, in that they are applicable to any type of censored discrete response, regardless of whether the data-generating process is defined by an intrinsically discrete process or by the rounding/grouping of continuous event times. Furthermore, the presented

methods are applicable in basically any field of research, as for example, in the social sciences, biostatistics, epidemiology, and many more. The U.S. unemployment data considered in this article is therefore only one of many possible examples. Further applications are presented in Tutz and Schmid (2016).

It is important to realize that all models considered in this tutorial can be fitted easily by use of standard software for binary regression modeling. The most important functions in `R` are `glm()`, `gam()` (of the **mgcv** package), and `rpart()` (of the eponymous package). In addition to the P-spline and CART methodologies considered here, many other options for semiparametric discrete time-to-event modeling exist in `R`. For example, **mgcv** provides a variety of alternative spline modeling tools such as cardinal splines and smoothing splines, which can be used for discrete hazard modeling by specifying the `bs` argument in `gam()` accordingly. Similarly, there is an alternative tree modeling approach developed by Bou-Hamad et al. (2009) that operates directly on the non-augmented time-to-event data. This procedure is implemented in the `R` package **DStree** (Mayer et al., 2014).

The basic functionalities required for applying the aforementioned software packages are all implemented in the **discSurv** package. Next to the functions used in this tutorial, **discSurv** provides additional functions to calculate, for example, measures for model evaluation like the concordance index (Schmid et al., 2017), and alternative tools for residual analysis.

We finally note that there exist a number of additional modeling options that are beyond the scope of this tutorial. These include, among many others, (i) regularized estimation via penalized optimization of the log-likelihood, which is useful for variable selection in higher-dimensional settings, (ii) random-effects and finite mixture model-

ing, which account for unobserved heterogeneity in the data, and (iii) competing-risks models, which extend the models considered in this tutorial by allowing for more than one target event. For details on further methodology, including semiparametric extensions, see Tutz and Schmid (2016).

Acknowledgements

The work of MS was supported by the German Research Foundation (DFG), grant SCHM 2966/1-2.

References

Agresti, A. (2013). *Categorical data analysis (3rd ed.)*. John Wiley & Sons, New York.

Bou-Hamad, I., Larocque, D., Ben-Ameur, H., Mâsse, L. C., Vitaro, F., and Tremblay, R. E. (2009). Discrete-time survival trees. *Canadian Journal of Statistics*, **37**(1), 17–32.

Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, J. C. (1984). *Classification and regression trees*. Wadsworth, Monterey.

Breiman, L. (1996). Technical note: Some properties of splitting criteria. *Machine Learning*, **24**(1), 41–47.

Croissant, Y. (2016). *Ecdat: Data sets for econometrics*. URL http://CRAN.R-project.org/package=Ecdat. R package version 0.3-1.

De Boor, C. (1978). *A practical guide to splines*. Springer, New York.

Eilers, P. H. and Marx, B. D. (1996). Flexible smoothing with B-splines and penalties. *Statistical science*, **11**(2), 89–102.

Ferri, C., Flach, P. A., and Hernández-Orallo, J. (2003). Improving the AUC of probabilistic estimation trees. In *European Conference on Machine Learning*, pages 121–132, Berlin Heidelberg, 2003. Springer.

Hastie, T., Tibshirani, R., and Friedman, J. H. (2009). *The elements of statistical learning (2nd ed.)*. Springer, New York.

Kalbfleisch, J. and Prentice, R. (2002). *The survival analysis of failure time data (2nd ed.)*. Wiley Inter-Science, New Jersey.

Klein, J. and Moeschberger, M. (2003). *Survival analysis: statistical methods for censored and truncated data*. Springer, New York.

Mayer, P., Larocque, D., and Schmid, M. (2014). *DStree: Recursive partitioning for discrete-time survival trees*. URL https://CRAN.R-project.org/package=DStree. R package version 1.0.

R Core Team (2017). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.

Schmid, M., Tutz, G., and Welchowski, T. (2017). Discrimination measures for discrete time-to-event predictions. *Econometrics and Statistics, accepted for publication*.

Schmid, M., Küchenhoff, H., Hoerauf, A., and Tutz, G. (2016). A survival tree method for the analysis of discrete event times in clinical and epidemiological studies. *Statistics in Medicine*, **35**(5), 734–751.

Therneau, T., Atkinson, B., and Ripley, B. (2015). *rpart: Recursive partitioning and regression trees*. URL https://CRAN.R-project.org/package=rpart. R package version 4.1-10.

Tutz, G. and Schmid, M. (2016). *Modeling discrete time-to-event data*. Springer, New York.

Welchowski, T. and Schmid, M. (2017). *discSurv: Discrete time survival analysis*. URL http://CRAN.R-project.org/package=discSurv. R package version 1.1.7.

Willett, J. B. and Singer, J. D. (1993). Investigating onset, cessation, relapse, and recovery: Why you should, and how you can, use discrete-time survival analysis to examine event occurrence. *Journal of Consulting and Clinical Psychology*, **61**(6), 952–965.

Wood, S. (2006). *Generalized additive models: An introduction with R*. CRC press.

Wood, S. (2017). *mgcv: Mixed GAM computation vehicle with GCV/AIC/REML smoothness estimation*. URL https://CRAN.R-project.org/package=mgcv. R package version 1.8-15.

Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **73**(1), 3–36.